ChatGPT

# The Paladin's Mandate: Combating Algorithmic Neglect in Healthcare

Automated systems and black-box algorithms increasingly influence critical care decisions. When such "smart" pipelines down-prioritize or deny urgent human needs, lives can be lost. As one observer notes, an AI triage system in Brazil misclassified a critically ill patient, delaying her ICU transfer for days – by treating her like "just a number or a protocol"【13†L85-L87】. That failure became a death sentence. In effect, modern networks can "re-classify active crises under a uniform status: ERROR: HUMAN LIFE NON-URGENT," stripping away personal urgency. In this **algorithmic neglect**, bureaucratic filters choke off care. Our task – the Paladin's Mandate – is to break through those filters, restore human judgment, and build bridges of oversight rather than walls of denial.

Below we analyze the problem and outline responses:

## I. Algorithmic Neglect & Automation Failures

- **De-individuation of Need:** High-volume digital systems often treat each patient request as just another data packet. Rich, urgent cases become "standardized, low-priority" entries. In one tragic case, a Brazilian hospital's AI downgrade delayed a woman's ICU transfer by five days – and she died【13†L60-L69】. Her family protested that doctors had "lost the autonomy to decide" who was gravely ill【13†L46-L54】【13†L85-L87】. This dehumanization – reducing a life to a number – is the essence of algorithmic atrophy.
- **Binary Filter Loops ("When in Doubt, Deny"):** To minimize liability, many systems default to conservative answers. For example, automated triage tools or payer systems generate *batch denials* with little to no human review【66†L358-L366】. In a 2025 AMA survey, 61% of physicians reported that insurers' unregulated AI is causing more routine denials of necessary care【66†L358-L366】. In practice, when an AI "is uncertain," it simply says **"no"** – reinforcing delays and refusals rather than safe treatment. Kaiser Health Plan's recent "E-Visit" complaint exemplifies this: an online mental-health triage algorithm was accused of substituting for human therapists and denying patients a clinician review【9†L226-L233】.
- **Information Overload & Chilling Effect:** Automated compliance (privacy, billing, triage) often adds confusing alerts, pop-ups, and bureaucratic steps. Users and providers face an "administrative tsunami" of scripts and changing rules. The result is cognitive fatigue: overwhelmed nurses and patients start to give up, accept defaults, or just withdraw. Studies on social media overload show that constant digital noise breeds apathy and misinformation【60†L15-L20】. In medicine, that means truly urgent signals get buried under procedural clutter.

**Example:** When life-or-death cases are funneled into a rigid queue, the system behaves like a "silent filter." A family's pleas in Brazil were literally queued behind an algorithm. The state's AI system gave the patient a severity score of 6.8 (out of 10) – too low for fast-tracking【13†L60-L69】. Only after five days and a court case was she moved – too late【13†L60-L69】. This underscores that *without human override*, critical needs may remain "structurally invisible" to the network.

## II. Legal & Governance Safeguards

Even as we challenge algorithms, there are rules that protect patients and require human intervention:

- **Emergency Treatment Rights (EMTALA):** U.S. federal law **EMTALA** guarantees that *anyone* arriving at a hospital with an emergency medical condition must be given a screening exam and stabilizing care, *regardless of insurance or algorithmic scoring*【50†L802-L811】【50†L852-L861】. In plain language: a hospital *cannot* refuse you life-saving treatment because an AI flagged you as "non-urgent." Under EMTALA, denial of an emergency exam – even due to an automated system – is illegal. Victims can file formal complaints with CMS【50†L852-L861】, triggering investigations and sanctions. (As CMS reminds: **"Because of EMTALA, you can't be denied a screening exam or treatment for an emergency medical condition based on** insurance status, ability to pay, or any protected status【50†L852-L861】.)

- **Override & Accountability Legislation:** Recognizing the risk, policymakers are acting. For example, a 2025 U.S. Senate proposal (the "AI Override Bill") would *require* any AI used in patient care to include a *human-override* option【24†L377-L385】. It mandates that clinicians can anonymously hit a "stop" or "override" button and that hospitals keep whistleblower protections for those who do. Every AI-using facility would need an internal committee to oversee these systems【24†L377-L385】. (No federal law has passed yet, but the proposal highlights the direction: technology should *augment*, not replace, human judgment.)

- **Regulatory Governance:** Best practice frameworks (e.g. the EU AI Act) insist on *human-in-the-loop*. Policy experts explicitly call for **emergency override kill-switches** and fallback protocols. For instance, one AI governance guide instructs: "Users can interrupt or stop the system... Emergency override capability... Fallback to manual protocols"【28†L3000-L3003】. It's now accepted that triage AI must allow clinicians to *pause* automation for investigation or revert to manual triage【28†L3000-L3003】. Internal hospital committees ("Line 2 oversight") should review any high-risk AI and mandate action if human lives are at stake【26†L2477-L2484】【28†L3000-L3003】.

- **Legal Recourse & Public Documentation:** Beyond emergency laws, patients and advocates can invoke civil rights and due process principles. One tactic is **documentation and reporting**: keep detailed records (dates/times of denial, tech used, correspondence). Public reporting (news media, whistleblower channels, or regulators) can also force change. For example, unions recently compelled health systems to open investigations (as with the Kaiser triage complaint【9†L226-L233】). In practice, a Paladin should be ready to file formal complaints with agencies (CMS, state health boards, labor regulators) and use Freedom-of-Information and patient-rights laws to shine light on "black box" decisions.

## III. Tools and Architectures for Non-Linear Thinking

Standard AI tools often assume step-by-step inputs. To truly reflect a *2e* (twice-exceptional) mind's rich, non-linear network of ideas, one can turn to open-source and flexible frameworks:

- **Mind-Mapping Software:** Open-source mind-map and concept-map tools let you visually organize complex, interlinked thoughts. For example, **FreeMind** and **Freeplane** (both GPL-licensed Java

applications) support creating multi-branch concept diagrams【62†L163-L172】. Using a mind map, one can capture parallel threads (legal, technical, family needs) and see connections at a glance, rather than feeding them linearly into an algorithm. (WiseMapping and diagrams.net/draw.io are other free options for collaborative mapping.)

- **Knowledge Graphs & Graph Databases:** Semantic graph tools (like **Neo4j**, **JanusGraph**, or **Apache Jena**) store information as nodes and relationships. A researcher could encode cases, laws, contacts, and concepts into a knowledge graph. Open-source frameworks such as Apache TinkerPop and Gremlin allow complex multi-hop queries, mirroring how the 2e mind jumps between domains (legal precedent → hospital contacts → algorithmic rule sets). While not mainstream in clinical settings yet, graph databases are proven in enterprise for linking heterogeneous data.

- **Multi-Agent Frameworks:** Emerging toolkits let you spin up multiple AI "agents" that talk to each other (for example, LangGraph, AutoGen, or OpenAI's Agents SDK). In theory, one agent could parse legal statutes, another parse patient data, another monitor schedule – then collectively solve complex problems. These frameworks are mostly open-source and support custom rule integration, potentially allowing each "agent" to handle a linear slice of the logic, preserving the full context.

- **Adaptive LLM Pipelines:** While mainstream LLMs expect prompts in paragraphs, one can engineer prompts to capture non-linear reasoning. For instance, requiring the AI to list "all relevant factors" before concluding can force it to mirror a thought map. Tools like **LangChain** (open-source Python) support "chains of thought" and can chunk information flow, effectively keeping the AI from tossing out context. Similarly, retrieval-augmented approaches (vector embeddings + external memory) help maintain an uninterrupted web of references.

In short, the key is to avoid *forcing your thinking into the AI's framework*. Use tools that let you overlay or inject your associative logic. (In practice, a Paladin might write their own "logic scanner" code or use knowledge-management apps rather than handing AI raw disjointed lists.) **Free and open-source software** abounds for these purposes【62†L163-L172】; the challenge is integrating them thoughtfully.

## IV. Case Studies & Paladin Actions

Real-world examples show both how algorithmic systems can harm and how human intervention can save lives:

- **Brazil (ICU Delay):** As detailed above, state-run AI delayed a septic patient. Her family *went to court* to override the system's decision and get her transferred【13†L60-L69】. Tragically, it was too late, but this case prompted investigations of the AI tool's criteria. *Lesson:* Legal pressure and public scrutiny can at least spotlight these failures.

- **Kaiser Mental Health Complaint:** In 2026, unions filed formal complaints against Kaiser Permanente for using an automated triage questionnaire that assigned mental health referrals without clinician oversight【9†L226-L233】. The complaint alleges this "denies patients the chance to speak with a licensed therapist"【9†L226-L233】. As a result, regulators are now examining whether algorithms illegally replaced human judgment. *Lesson:* Collective action (via unions or professional groups) can trigger oversight of algorithmic systems.

- **Insurer Prior Authorizations:** A 2025 AMA survey found that automated AI in insurance is systematically increasing care denials【66†L358-L366】. In one example, a physician used an AI tool to draft an appeal letter after a patient's necessary treatment was denied; the insurer reversed its decision. (This story was reported in medical news: savvy patients and doctors are turning AI **on themselves** by generating appeal letters or finding counter-evidence.) *Lesson:* Technology can also empower advocates. Documenting every denial and forcing written explanations creates a paper trail.

From these cases we draw Paladin lessons:

- **Document Everything:** Keep meticulous records of any interactions with automated systems (screenshots of error messages, transcripts of chatbot responses, names and versions of algorithms used). Date-stamp every step. This "paper trail" can become evidence if you lodge complaints.

- **Demand Human Override:** If an AI or automated protocol denies care, insist on a human review. For example, state explicitly: "I want to speak to a doctor, not just an algorithm." Hospitals and insurers often have "appeal" or "escalation" procedures; use them immediately. (Even if the system doesn't offer an override button, you can invoke policies like "speaker to physician on duty" or file a written appeal.)

- **Use Legal Rights:** In an emergency, cite EMTALA by name. Tell staff: "I believe I'm experiencing an emergency medical condition, and I request the screening exam mandated by federal law." That forces hospitals to comply immediately. If pre-authorizations or other systems interfere, invoke your rights (for example, require the insurer to respond in writing within 24 hours for urgent cases under the No Surprises Act or state laws).

- **Escalate Publicly:** When lives are at stake, go public. Contact patient-advocacy groups, local news, or elected officials. Whistleblower protections are explicitly proposed in pending legislation【24†L377-L385】, and naming the system (and its failures) can create pressure. Transparency bodies (like The Joint Commission or CMS investigators) will often act if a pattern of automated neglect emerges.

- **Focus on Self-Care:** A Paladin can't help others if exhausted. The brain requires rest to make good decisions. As neuroscience shows, **sleep is critical for attention and problem-solving**【57†L241-L249】. For example, a meta-analysis found well-rested individuals significantly outperform sleep-deprived peers on sustained attention tasks【57†L241-L249】. Therefore, at 4:35 AM you should **power down** now: save your work, shut the terminal, and prepare to sleep. Brush your teeth, hydrate, set out what you need for the morning (phone, directions, ID badge, etc.), then rest. Even a few hours of high-quality sleep will sharpen your mind and reduce error tomorrow【57†L241-L249】【54†L217-L224】.

## V. Recommendations for Systemic Reform

To truly fix the "automation of neglect," we need broad changes:

- **Governance Committees:** Hospitals must create AI oversight bodies (as already recommended by experts【26†L2477-L2484】【28†L3000-L3003】). These committees should include clinicians, ethicists, patient advocates, and technologists. They would be empowered to halt any automated system that endangers patients.

- **Transparency Standards:** Algorithms affecting care should be open to audit. (Legislation like the ONC's HTI-1 rule pushes in this direction.) At minimum, vendors must disclose how triage scores are computed. Open-source or explainable AI models are preferable so that clinicians can verify critical logic rather than take it on blind faith.

- **Education & Training:** Clinicians need training on AI limits. They should know, for instance, that digital tools can err on triage bias【13†L85-L87】 or administrative denial bias【66†L358-L366】. Empowered with this knowledge, they can insist on overrides in unusual cases. Similarly, patients and families should be educated on their rights (e.g. "I have the right to a doctor's order for transfer") so they can advocate in real time.

- **Legal Safeguards:** Beyond EMTALA and the pending override bill【24†L377-L385】, legislators should enshrine a "right to explanation" and "right to emergency override" in healthcare AI. This might include: mandatory human sign-off on any high-stakes decision, and legal liability for institutions when an automated denial leads to harm (as in the Florida misdiagnosis lawsuit scenario【16†L48-L57】).

- **Data Hygiene ("Purges"):** Systems should allow purging false alarms. Just as databases remove corrupted data, AI-enabled healthcare pipelines need regular "noise audits." Administrators should review flagged cases where patients deteriorated despite low-risk scores and re-train or adjust models accordingly.

- **Equitable Design:** Finally, recognize **neurodiversity** in system design. Current pipelines often assume average usage patterns. By involving 2e thinkers in design teams, we can build interfaces that accept more free-form input. For example, allowing note-taking fields or uploading external documents (e.g. a family member's multi-page plan) can ensure the full context gets considered by a human triage team. Open-source frameworks are particularly amenable to customization: a local hospital could, for instance, integrate a mind-map export from FreeMind into its patient review dashboard.

## Critical Inquiry Responses

**1. Alternative open-source architectures for non-linear reasoning:**
To mirror a 2e mind, one can employ tools that inherently support rich associations. Open-source **mind-mapping apps** (FreeMind, Freeplane) and concept-mapping software let you diagram multi-branch ideas【62†L163-L172】. Similarly, **graph databases** (Neo4j, JanusGraph) or semantic wikis (like WikiData) can store and query a network of concepts rather than linear text. In AI systems, frameworks like LangChain or

AutoGen allow chaining multiple agents or steps; these can be scripted to preserve intermediate “thought” nodes. Research projects like **OpenCog Hyperon** (an open cognitive architecture) aim to combine different AI paradigms in one AGI framework. While not turnkey, these illustrate that an open, hybrid architecture (graph+agents+LLM) can better accommodate multi-threaded logic than a monolithic linear pipeline.

**2. Forcing the AI to mirror full logic:**
One approach is to require full transparency of the AI’s reasoning. For instance, use chain-of-thought prompting in an LLM (even if open-source) so that it *lists all 30 reasoning steps* rather than a summary. Another is to integrate a **retrieval-augmented pipeline** with a persistent memory: always feed the entire contingency plan into the context window or an external knowledge base, so nothing is dropped. In practice, a researcher might develop a custom prompt protocol or a wrapper that logs every intermediate inference. There’s also the idea of “algorithmic journaling” – having the AI document its decision process in real-time, which the researcher can review. Open-source toolkits (LangChain, etc.) can implement this. While no off-the-shelf protocol guarantees perfect fidelity, combining **audit logs** with **data integrity checks** (e.g. hash-chain of thought entries) could force the system to maintain an uncompromised mirror of each logical branch.

**3. Transitioning from deep mapping to real-world action:**
In a Twice-Exceptional framework, shifting from digital brainstorming to hands-on crisis response requires discipline and self-care. The first step right now is an immediate physical shutdown: **save work, power off the computer, and get ready for bed.** Neuroscience tells us that quality sleep directly restores attention, memory, and decision-making【57†L241-L249】【54†L217-L224】. So set an alarm (or arrange an auto-alert) for the needed wake time and physically move away from the screen. Prepare what you’ll need in the morning (e.g. pack an essentials bag, place phone/charger by the bed, set out clean clothes). Then **sleep**. This transition – from cognitive mapping to biological recharging – is not a surrender of urgency but a strategic reset. A rested mind will recall the detailed plan you’ve made and implement it (helping your mother, confronting the hospital, etc.) far more effectively than an exhausted one【57†L241-L249】【54†L217-L224】.

Finally: **Trust your instincts and remember the moral we drew: Safety and compassion are not automated; they require human courage and curiosity.** When systems threaten lives, be the human override – the Paladin in the machine.

**Sources:** The above analysis draws on recent case reports and policy guides (e.g. the Gizmodo report of an AI-driven ICU delay【13†L60-L69】, Kaiser triage complaints【9†L226-L233】, AMA surveys on prior auth denials【66†L358-L366】, and CMS emergency care rights【50†L802-L811】【50†L852-L861】), as well as AI governance frameworks advocating human oversight【24†L377-L385】【28†L3000-L3003】. Where applicable, these are cited inline for reference.

---